*THE PSYCHOLOGICAL HEURISTICS OF LEARNING**

By Ernest R. Hilgard

STANFORD UNIVERSITY, STANFORD, CALIFORNIA

Our discussion centers upon the teaching-learning process rather than upon the finer structure of what goes on in the brain when we learn; hence the title "heuristics" of learning. In the present-day language of science we are somewhat out of the basic science area into the applied science portion of the spectrum, although, as we all know, these distinctions are never sharp.

Until recently, the payoff from the basic study of learning to the applied areas has been quite limited, to the embarrassment of some of us who have worked for many years on laboratory studies of learning. The topic has in fact been one of high prestige within the experimental laboratory in psychology, perhaps the favorite topic for laboratory study in psychology over the last 30 years. Yet when it comes to teaching reading, writing, or arithmetic, the advances owing to this enormous investment in the science of learning have had but slight consequences. We know a great deal about how a white rat learns a maze, but when we teach a boy to ride a bicycle, we usually give him the bicycle and let him teach himself, without worrying about our carefully studied principles of task analysis, distributed practice, or prompt reinforcement. As we now turn to the heuristics of learning we must ask ourselves how we can help the boy with his bicycle rather than how we explain the rat's threading its way through the maze.

If I call attention first to the failures and limitations of the psychology of learning, I do not wish to give the impression that all is lost, for I believe that some of the contemporary developments are very promising. But first let me comment on some of the limitations. The first limitation is remoteness from actual, concrete learning situations that we are called to manage; a monkey in a cage is not a student in a classroom. But the limitations are found as well in studies of high pertinence to the classroom: relevance alone does not guarantee a sound psychology of learning. Let me illustrate from some studies of college teaching.

Many of us are college or university teachers and most of us have participated from time to time in discussions on what makes a good teacher and on what arrangements are best for instruction. Higher education is big business, and there is no reason why we should not try to introduce greater efficiency into it through some sort of cost-benefit analysis of different kinds of teaching. One obvious candidate for study is class size.

Despite the fact that the teachers remembered by many of us as providing our great experiences as undergraduates were the brilliant lecturers, there is a strong feeling that the teaching advantage lies with the small class. The preference for the small class is shared by students. A recent study at Stanford of the seniors graduating in the class of 1968 yielded the data of Table 1, which assigns rewarding experiences in the order: independent work, seminars, medium-sized lecture-discussion, and large lectures.

What of the efforts to make class size a topic for research? A careful analysis of nearly 100 studies of class size in college over the past 40 years leads inescapably to one conclusion: There are no demonstrable differences in results (as judged by final examinations) of small classes versus large ones, or even individual study whether with a tutor or individual study totally unsupervised.[4] If we take these bits of evidence at their face value, the efficient thing to do in college or university teaching would be to prefer two methods. The first is to maintain the very large lecture, with, of course, the best lecturers available; the other method would permit students to work very largely alone. Thus we would provide for a maximum utilization of staff time. These conclusions are shocking and, as a matter of fact, most of those who conducted the investigations under review were unwilling to accept the results of their own investigations. Some of the limitations which they note are the following:

(1) Most of the investigations have been concerned with introductory courses taken by large numbers of students. While a great deal of new information is acquired in these courses, they probably do not provide the most satisfying experiences for students in any case. The satisfactions from independent work and small seminars reported by the Stanford seniors (Table 1) were doubtless mostly in advanced work in the areas of the major.

(2) In order to carry out a study comparing two methods, it is necessary to have some common criterion by which to judge progress. This is usually a final examination based on either a textbook or a list of readings; hence it can test only what the methods have in common, not anything unique to one method or the other. (Of course novel problems may be set in the examination, but these may test general intellectual level and not reflect differences in teaching.) The common intellectual content represented by a reading list can be mastered as well by independent reading as by lecture or discussion.

(3) More subtle objectives of teaching, such as generating enthusiasm for intellectual inquiry, gaining in self-understanding, coming to identify with a profession, changing social values or attitudes, are not reflected very well in most of the studies.

My purpose in presenting this material is not to revolutionize college teaching, or to propose substitutes for examinations as we now conduct them, but to call attention to a failure of research if we are unwilling to be guided by it. I fear that the demonstration has not been made soundly enough for us to be convinced, although the studies were carefully done because those who did them knew that the stakes were large, and each one felt that he had improved upon the methods of the earlier studies.

TABLE 1. *Of those reporting high experience, percentage evaluating experience as highly rewarding, by classroom situation, by area of major (Stanford seniors, Spring 1968).**

	Engineering (%)	Science (%)	Social Science (%)	Humanities (%)
Independent work	86	93	85	79
Seminars	89	80	83	87
Medium-sized lecture discussion	54	46	45	54
Large lectures	7	18	17	23

* In *The Study of Education at Stanford: II. Undergraduate Education* (Stanford, Calif.: Stanford University, 1968), p. 59.

The same kind of negative result holds for studies of elementary education. There are really no consistent differences to be found between teaching reading by the whole-word method or by the phonetic method.[3] We really do not know whether there would be an advantage in using a different initial alphabet in English with beginners, or whether we should start teaching reading earlier than we now do, or whether we should postpone it until later, as they do in Scandinavian countries and in Russia, apparently with success. These are practically important choices which hang upon knowledge, and yet our knowledge is too insecure to permit clear formulations of desirable practice. Despite all the attention that has been given to child development and to sex differences, we are unable to account satisfactorily for a ratio of some five boys to one girl in our remedial reading classes.

Somewhere along the line, psychologists, along with educators, have failed to come up with a truly *responsible* applied psychology of learning. Let me summarize the state of the art as it existed perhaps a decade ago:

(1) There were thousands of experimental investigations of reading, but they had not led to agreement on the preferred methods of teaching.

(2) There were upwards of a hundred quantitative studies of class size in college teaching, with the verdict that one method was no better than another, but the verdict was not accepted. Lack of relevance is not the obstacle in studies of reading or class size.

(3) There were thousands of laboratory studies in the name of a basic science of learning, on conditioned responses, on motor skills learning, on nonsense syllable learning, with animals and human subjects, largely irrelevant to the solution of practical problems, or at least lacking in any capitalizing on their potentials of relevance through the bridging experiments to make relevance explicit.

The obvious need was not for more of the same, but for something different. There is no reason to expect new studies of the old kind to lead to anything more definitive than the old studies. The temptation is to continue—which is, I suppose, a common disease in what Kuhn[5] has called "normal" science.

Some efforts were indeed made over the years to break out a little from the standard patterns. As the motion picture became cheaper and easier for the teacher to project, educators hoped that visual aids would provide new dimensions to teaching, and then the tape recorder added the audio dimension, so we had audiovisual aids (Brown and Thornton[2]). Countless studies of these led to the same old conclusion: One method is as good as another.[8] Yes, people can learn from films, perhaps a little better than from a very poor instructor, but no better than from an average instructor. The hopeful thing in all of this is that when people gather together (or work alone), *if* they want to learn, given some learning materials, they can be shown to learn. The only problem is one of efficiency, and through the years notebooks, workbooks, laboratories, films, tapes, lectures, discussions, textbooks have all helped people to learn, but never with any dramatic changes owing to a new technology.

Two new hopeful processes have come along which may indeed break this logjam. The first of these is *programmed learning* in general, and the second is *computer-assisted instruction,* a special derivative and extension of programmed instruction.

Although there had been earlier teaching machines in the 1920's,[6,7] programmed learning took off from the work of B. F. Skinner in the 1950's.[9,10] He had done authoritative work on what has come to be called *operant conditioning*, chiefly with rats and pigeons, but developed a few simple principles that could be applied to any kind of training procedure. He and his students have turned out to be remarkably effective applied psychologists of learning, despite the basic science attitudes inculcated during some 30 years of precise studies of animal learning in the laboratory. Here, then, is the kind of payoff that a science of learning might have wished for.

The applications have extended to animal training, a curious lack on the part of some who have through the years worked on animal learning. It is Skinner's students who train the dolphins and other performers in the various marine Worlds which are now so popular. His methods are used in drug testing in pharmaceutical houses, in psychotherapy with autistic children and with schizophrenics, and in many other areas of application outside the schools. The advantage of his particular kind of formulation is that it tells you what to look for and what to do, and these are the marks of a science on the way to becoming a technology.

Let me summarize the Skinner system of operant conditioning to indicate what I mean by its technological simplicity.

First, the learner always comes to a given learning problem with something he can already do. Thus, he may know how to count before he tries to learn how to add or subtract, he knows how to talk before he learns how to read, and so on. This is described as the *operant level* at the time a new task is undertaken. Operant level is, in fact, a very complex matter of prior training, memory, individual differences, and motivation, but it reveals itself by the responses that the learner makes when he begins a new task, from the viewpoint it is just a matter of beginning to teach on the basis of what the learner already knows and can do.

Second, because of this operant level the learner does *something* in the presence of the new task. He characteristically varies his responses somewhat, in accordance with what traditionally has been called trial-and-error. In any case, when he does something that approximates a desired performance he is given some sort of immediate reward or *reinforcement*, as it is called within this system. A reinforcement is anything which tends to increase the probability that when next exposed to the same opportunities for response he will tend to do what he last did; the reinforcement may come by way of a chocolate drop, a pat on the back, or a verbal OK.

Third, absence of reinforcement leads to *extinction*, so that if behavior is not reinforced, its probability of recurring will be reduced, thus giving the opportunity for behavior to vary and for more appropriate behavior to appear and to be reinforced.

Fourth, by the skilled use of selective reinforcement and extinction, behavior can be made to move from a crude approximation to a more refined and acceptable performance. This process of directing the behavior in desirable directions is called *shaping*, and represents the essence of the new technology.

A trainer or a teacher who knows only about operant level, reinforcement, extinction, and their appropriate patterning in shaping, is ready to roll up his sleeves and go ahead. There are subtleties within the shaping process, such as the timing of reinforcements, the use of various schedules of reinforcement, and so on, but these are accessory principles, like learning how to tune the carburetor after you know how an internal combustion engine works.

While giving full credit to Skinner and his followers for the applied consequences both of his theory and of his inventiveness, let me point out that the theoretical support for his technology can come from sources other than his own theory. For example, the cognitive theorist, who has more respect for intellectual understanding than the stimulus-response psychologists among whom Skinner is included, is impatient with a theory of learning that limits itself to talking about small steps (as in a program), or responding, or reinforcement. He rightly points out that any subject matter has some kind of organization within itself that a learner must comprehend or understand if he has really mastered it. The cognitive theorist thus looks for the effects of the organization or structure of knowledge upon the ease with which the learner acquires what he needs to know.

Because the program constructor is likely to be talking the language of his technology, he may well fail to communicate all that he himself believes. Thus in stressing the reinforcement of responses he may in fact be neglecting to say anything about what is really being learned.

The little responses that fill in the blanks at the end of a program, or the words that the student points out on the television screen with his electronic pencil, are *not* what is being learned, although they may be *indicators* of what is learned. Suppose that in learning to extract the square root of 25 you get the answer 5, and write it down. The "5" then gets reinforced, because it is correct. Did you learn the response "5," or did you learn to extract the square root? When a rat runs a maze, and gets to the end-box, and eats the food that serves as a reinforcement, is he learning to eat? Obviously the response at the end is merely a special output that shows that the essential responses along the way have been made, or, in cognitive terms, that the essential relationships have been understood. A program could be written that would have all the answers either the word "right" or the word "wrong," as in a true-false examination. Obviously more would be learned than to write the words "right" and "wrong" in the spaces provided. The point being made is that cognitive learning can be taking place under the technological arrangements of operant conditioning or programming. We are easily misled into believing that the whole process is far more mechanical than it is, merely because it makes use of mechanical aids.

I have moved into a discussion of programmed learning without specifying very clearly what it is. In its earlier form, the program progressed in small steps, taking the learner from where he is at the beginning to where the teacher wants him to go. The early teaching machines and programmed books incorporated such a procedure; these came to be known as linear programs. Another kind of program developed very shortly, and it is really out of this second kind of program that the computer-assisted instruction evolved. The was called a branching program. All learners do not follow the same path through the program, but the next steps are contingent upon earlier ones and are sometimes based upon the learner's own preferences. The computer can facilitate such programs because it provides maximum flexibility, will do what it is told, and does not forget its instructions.

An advantage of programmed learning, and of the teaching machine, is that there is a record of the student's progress, of errors made, of amount learned per unit time. This is in some respects the most significant advance over ordinary teaching methods. Most teachers really do not know what the learners are doing; they trust to a student's occasional smile or a nod of the head for assurance that the student is listening. The questions that come up later in the hour often as not show that the student was not listening, that he really didn't hear what had gone before. When examination time comes, teachers are often disappointed because of many students' learning failures, and pleased with what the brighter ones know—but they have little idea what their own teaching had to do with it. In programmed learning you know where the student is and what he is doing; if he progresses through the program, both he and the teacher have the satisfaction of knowing that he learned from it. Computer-assisted instruction not only has this same advantage of keeping a record, but it has the further advantage that the computer can make *computations* so that an analyzed record is available for each student at any time.

The potential contribution of the records kept by the computer is enormous. No matter how elegant our mathematical models may be, there is necessarily a large element of empiricism that enters into the teaching of any given subject matter to a child of a given age and experience. One of the errors of laboratory studies of learning in the past was the neglect of care in discrimination of the content learned, so that we trusted far too much to laboratory devices which were convenient to illuminate partial processes, but which told us nothing about the parameters important in learning geometry, a foreign language, or art appreciation. The same principles may be involved, but these principles will not prepare instructional materials appropriate both to the subject matter and the learner. But with the aid of the computer a vast amount of empirical data can be collected and summarized, with a view to optimizing learning. It is not enough that learning should take place—it should take place as efficiently as possible. I believe that our basic science of learning will be better as a consequence of these large-scale efforts at optimizing learning. Never before have we had a learning laboratory with so much processed data.

Having indicated my optimism with respect to computer-assisted learning, I wish to address the rest of my remarks to the problem of the proper place that I see for it in the total educational process, and what I see as its limitations.

Let me first acknowledge the promise that I believe such instruction holds for the efficient teaching of all manner of skills, information, and appreciation. I see no inherent limitation to subject matter with fixed answers, such as mathematics, grammar, map reading, historical chronology, or foreign language vocabulary. It is possible to teach poetry or creativity as well, and programs now exist to serve these purposes.

Since I have indicated optimism with respect to these new developments, I want to make a few remarks about keeping computer-based instruction in its place. I think, for example, that there are two things that ought to be said very pointedly about computer-based instructions that are themselves limitations. One is that it is sedentary, since you are pretty well tied to a terminal when you're working with a computer, whereas a good deal of learning is actually on the hoof, moving around, as in conversation with people whose rejoinders are more capricious than those of the computer. If we are to encourage the spirit of inquiry, we have to poke around in the laboratory or in the shop or in the library or go on field trips, and this poking around is not possible and must not be neglected in a warm-hearted program of instruction. The second thing, which is related to the fact that it is sedentary, is that it is too individual. It's very nice to have a slogan of individualizing education, and there are great advantages in some individualization, but life isn't lived quite that way. We live within restraints. Any of us who work in larger laboratories know that even in basic science, we are restrained by the demands of our fellows and the expensive equipment that has to be shared. Life is inevitably social, and one has to be concerned with the social aspects along with the individualizations of learning, so that we really don't want to completely individualize education. We want a lot more than we have, but we don't want it all the way. Learning to participate in a play or to take part in an athletic team or to conduct a group discussion on a controversial issue—these are matters that require social learning.

The role of the teacher, as I see it, is to do these other things. The teacher will be freed by the computer from spending much time in the communication of knowledge and in the questioning and answering that now goes on so commonly in the classroom. When this is the case, the ordinary teacher will have to be trained to do what the best teacher now does.

The teacher can take responsibility to see that the student learns to *initiate* inquiry on his own. While the computer can provide a range of opportunities and can even engage in individual guidance, I doubt if it will ever do as well as a skilled teacher in fanning a faint spark into a glowing interest. Recent work on social learning theory[1] has shown that *imitation* is one of the neglected areas in the psychology of learning, and the imitation of a teacher as an adult model may have great influence upon what is learned.

The teacher can help the student to gain *a favorable image of himself* as a learner and as a creative person. While the reinforcements of the computer will help, direct social approbation is an even more powerful reinforcer. I was greatly impressed by something that happened many years ago when I was working with young chimpanzees, along with Professor Yerkes at the then Yale Laboratories of Primate Biology in Florida. So as not to introduce experimenter bias into the session, we were concealed behind a screen while the chimpanzee went about his puzzle-solving. He solved the problem, all right, and a banana appeared automatically as a welcome reinforcer. He picked up the banana, but then he sought out the screen and peered behind it to show *us* the banana and get our commendation before he sat down to eat it. The "computer" had delivered his reinforcement, but he wanted ours in person. I suspect children are like that, too.

The teacher also has a role in directing the student toward *effective participa-*

tion with others. While I am against making everybody into extroverts, human life is inescapably social, and an effective person has to learn to cooperate with others in solving problems, in making plans, or in carrying out a cooperative enterprise, whether at home, at school, at work, or in the community. The skills of social participation, of leadership and of followership, of tolerance of opposition and of frustration, of social conflict resolution, can be learned only through exercising them. The discriminations are too difficult, the response interchanges too rapid, for them to be well-programmed. Even after social skills and practices have been studied through a program, they have to be exercised or they will not persist.

What this amounts to, then, from what we know about how an individual learns, and how he can be aided by those who wish to aid his learning, is that computer-assisted instruction is soundly grounded in what we know about learning, although its usefulness does not arise exclusively from any one of the prevailing theories and it itself contributes to them; it will not make the teacher dispensable, but will alter the teacher's functions in such a way as to require the usual teacher to do what only the exceptional teacher now does well. This is itself an important opportunity for teacher-training institutions, as they prepare teachers for the schools of the future.

[1] Bandura, A., and R. H. Walters, *Social Learning and Personality Development* (New York: Holt, Rinehart, and Winston, 1963).

[2] Brown, J. W., and J. W. Thornton, Jr., eds., *New Media in Higher Education* (Washington, D. C.: National Education Association, 1963).

[3] Chall, Jeanne, *Learning to Read: The Great Debate* (New York: McGraw-Hill, 1967).

[4] Dubin, R., and T. C. Taveggia, *The Teaching-Learning Paradox: Comparative Analysis of College Teaching* (Eugene, Ore.: University of Oregon, 1969).

[5] Kuhn, T. S., *The Structure of Scientific Revolutions* (Chicago: Univ. of Chicago Press, 1962).

[6] Pressey, S. L., "A Simple Apparatus Which Gives Tests and Scores—and Teaches," *School and Society*, **23**, 373–376 (1926).

[7] Pressey, S. L., "A Machine for Automatic Teaching of Drill Material," *School and Society*, **25**, 549–552 (1927).

[8] Schramm, W., ed., *New Teaching Aids in the American Classroom* (Stanford, Calif.: Institute for Communication Research, 1960).

[9] Skinner, B. F., "The Science of Learning and the Art of Teaching," *Harvard Educ. Rev.*, **24**, 86–97 (1954).

[10] Skinner, B. F., "Teaching Machines," *Science*, **128**, 969–977 (1958).